techUK
FOR WHAT COMES NEXT

Mapping the Responsible AI Profession: Current Practice and Future Pathways

A Field in Formation

April 2025

About techUK

techUK is a membership organisation launched in 2013 to champion the technology sector and prepare and empower the UK for what comes next, delivering a better future for people, society, the economy and the planet. It is the UK's leading technology membership organisation, representing around 1,000 members across the country, including 700 SMEs. We are a network that enables our members to learn from each other and grow in a way which contributes to the country both socially and economically. By working collaboratively with government and others, we provide expert guidance and insight for our members and stakeholders about how to prepare for the future, anticipate change and realise the positive potential of technology in a fast-moving world.

About techUK's Digital Ethics Programme

In an increasingly digital world, it's important that technology is used to improve and enhance the quality of people's everyday lives. Embedding ethical principles, such as transparency, accountability and explainability, into the creation of products, tools and services is essential for building public trust and confidence in technology. techUK focuses on resolving some of the most difficult ethical challenges, to ensure tech works for people and responsible innovation can flourish.

Table of Contents

Executive Summary

As artificial intelligence (AI) becomes adopted across the UK economy, a new professional class has emerged as essential: Responsible AI (RAI) practitioners. These individuals ensure that AI systems are developed and deployed ethically, safely and fairly.

The UK's ambition to become a global AI innovator while maintaining strong ethical standards creates unique opportunities and challenges for RAI practitioners. Our analysis reveals that RAI practice stands at a critical juncture. The field has evolved from an emergent discipline into a field with an essential organisational function, albeit one still finding its formal structure and boundaries. The diversity of practitioner backgrounds has proven to be a strength, enabling organisations to address the multifaceted challenges of ethical AI implementation, while creating challenges for professionalisation efforts.

The growing complexity of AI systems demands increasingly sophisticated governance approaches. Organisations recognise that effective RAI practice requires both dedicated expertise and distributed responsibility, with practitioners often serving as orchestrators rather than the sole owners of AI ethics and governance. Without these professionals to operationalise ethical principles and support the assurance ecosystem, regulatory requirements become merely aspirational, rather than practical. Yet three critical gaps currently undermine responsible AI practitioners effectiveness and threaten the UK's AI leadership ambitions and the success of the assurance ecosystem.

Our mapping of the current landscape reveals the first critical gap: the absence of clear role definitions and organisational placement. Organisations lack consensus on where RAI functions should sit, what authority they should hold and how they should interact with development teams. This creates inconsistent implementation, dilutes effectiveness and makes accountability difficult to establish.

An analysis of practitioner profiles identifies the second gap: the lack of structured career pathways. Current practitioners come from remarkably diverse backgrounds - philosophy, computer science, law, and the social sciences - bringing valuable perspectives but creating challenges for standardisation. Career paths remain non-linear, similar to privacy practice before established professional routes emerged.

Our assessment of educational opportunities highlights the third gap: the absence of standardised skills and training frameworks. While academic programs and professional certifications are emerging, organisations have limited guidance on which

credentials signal genuine expertise. This hampers talent identification, development and retention.

These gaps create tangible business risks: inconsistent ethical implementation, potential regulatory non-compliance, reputation damage and barriers to establishing stakeholder trust. They also threaten the broader economy by potentially hindering the UK's ability to establish leadership in responsible AI innovation.

Forward looking organisations could establish RAI roles with clear mandates and direct reporting lines to leadership, while investing equally in technical capabilities and governance skills. Professional bodies could develop flexible certification frameworks that recognise multiple pathways to expertise while establishing distinct boundaries between the ethical, auditorial and compliance functions of RAI practice. Policymakers could support industry collaboration and invest in educational pathways that develop both technical and ethical competencies.

Just as privacy experts became indispensable during the internet’s expansion, responsible AI ethics practitioners are now becoming the essential human infrastructure for the UK's AI future. By addressing these gaps, the UK can cultivate user trust, demonstrate regulatory readiness and attract investment - building a foundation for adoption and confidence in AI. This paper provides a roadmap for cultivating the professional ecosystem necessary to ensure that AI development in the UK remains both innovative and aligned with our societal values and ethical standards.

Introduction

The rapid mainstreaming of AI has created a fundamental shift in how organisations approach AI governance and ethics. What was once primarily a theoretical discourse or auxiliary function has evolved into an urgent operational imperative, often seen on the board agenda, with organisations scrambling to establish robust frameworks for responsible AI (RAI) implementation. At the heart of this transformation lies a pressing question: who, precisely, is responsible for responsible AI?

This question has given rise to a new class of professionals – RAI practitioners – whose roles are as diverse as they are essential. RAI practitioners are generally tasked with interpreting and operationalising best practices for the ethical and safe design and deployment of AI systems[1]. These individuals and teams serve as bridges: connecting ethical principles to practical implementation, weighing technical capabilities with societal implications and aligning organisational innovation with regulatory compliance[2]. Their work requires the constant navigation of trade-offs while facilitating vital cross-team collaboration.

The UK government's current aim is to foster increased AI adoption and diffusion across the economy. Key to achieving this will be cultivating greater trust and confidence in AI systems and credibility in the professionals who safeguard them. This is why the role of the RAI practioners is crucial and the need to support the development of this profession is vital. However, we currently lack clear pathways for individuals to enter the responsible AI profession, creating uncertainty for hiring managers and impeding the development of a robust assurance ecosystem and supportive skills programmes as recommend in the recently published AI Opportunities Action Plan.[3]

---

[1] Rismani, Shalaleh, and AJung Moon. "What does it mean to be a responsible AI practitioner: An ontology of roles and Skills." *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 8 Aug. 2023, pp. 584–595, https://doi.org/10.1145/3600211.3604702.

[2] Cocchiaro, Mariangela Zoe, et al. *Who Is an AI Ethicist? An Empirical Study of Expertise, Skills, and Profiles to Build a Competency Framework*, 2024, https://doi.org/10.2139/ssrn.4891907.

[3] "AIOpportunities Action Plan." *GOV.UK, 2025*, www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan.

The landscape of AI governance is rapidly evolving[4]. The enforcement of the EU AI Act, coupled with emerging regulatory frameworks worldwide, has created demand for professionals who can navigate the technical and ethical dimensions of AI implementation. In response to this demand, organisations have established new roles focused on responsible AI practices. RAI practitioners are expected to take on more strategic positions within their organisations. This shift goes beyond technical oversight to encompass broader organisational change management and strategic decision-making around AI deployment.

The UK's principle-based, agile and context-specific regulatory approach to AI makes the role of responsible AI practitioners even more critical. RAI practitioners are often tasked with operationalising the five ethical principles underpinning the UK's pro-innovation approach to AI regulation: safety, transparency, fairness, accountability and contestability. Without these practitioners, these principles could remain aspirations rather than operational realities. These professionals help their organisations navigate regulatory requirements to maintain consumer, citizen and industry trust, while also helping companies maintain credibility. Organisations are no longer asking whether they need dedicated responsible AI expertise, but rather how to structure these roles effectively and what competencies to prioritise.

The career pathways leading to RAI practice (i.e. chief ethics officers, heads of AI ethics and responsible AI keads within organisations) are remarkably diverse, reflecting the field's multidisciplinary nature. Current RAI practitioners come from varied backgrounds including philosophy, compliance, computer science, law, the social sciences and business management. This diversity brings rich perspectives to RAI implementation and should be viewed as a strength. However, some have compared the current state of RAI practice to privacy practice 20 years ago, when defined career paths had not yet emerged. As the profession matures, more standardised educational and career pathways will develop, even though maintaining diversity in professional backgrounds will remain valuable.

While this diversity in professional backgrounds strengthens the RAI field, it also presents significant challenges for establishing RAI as a recognised profession. A key next step lies in differentiating between key functional streams within AI governance. Responsible AI practitioners - professionals who implement ethical, safe and fair AI

[4] Casovan, Ashley, et al. “AI Governance in Practice Report 2024.” *AI Governance in Practice Report 2024*, iapp.org/resources/article/ai-governance-in-practice-report/.

systems within organisations - must navigate increasingly blurred boundaries between related domains. There is a growing risk that ethics work may drift into pure compliance functions, potentially reducing nuanced ethical and sociotechnical deliberation to regulatory checklist exercises.

For RAI to mature as a profession as privacy practice has done, the boundaries between ethics, audit and compliance roles must be clearly defined while preserving the integrity and effectiveness of each function. Organisations currently recognise that academic AI ethics serves a different purpose from practical responsible AI implementation in business contexts, and that despite the two fields complementary nature, each requires a different skill set from its practitioners.

The business implications of addressing these professional development challenges are substantial. Without structured professionalisation, organisations may face the inconsistent implementation of ethical AI principles, the erosion of stakeholder trust and potential regulatory complications that could hinder innovation and competitive advantage. The absence of professional standards leaves companies vulnerable to reputational damage and creates barriers to international collaboration and commerce in AI systems.

Furthermore, the lack of clear career pathways makes it difficult for organisations to identify, attract and retain talent in critical RAI roles. These business risks extend beyond individual practitioners or organisations to potentially impact the broader economy and the UK's ability to establish leadership in responsible AI innovation. Professionalising the field has important specific benefits, but, also, numerous potential drawbacks that must be addressed.

Regardless of the structural approach ultimately adopted, the professionalisation of RAI requires both technical expertise and ethical understanding. The path forward must build on existing professional foundations while addressing the novel challenges that AI presents to society.

This paper maps the current state of the UK’s RAI profession and demonstrates why the role of RAI practitioners as the human infrastructure upon which effective AI governance and ethics ultimately depends is essential to the UK’s AI future. By understanding the current landscape, competencies and developmental needs of these professionals, we can more effectively target investments in capability-building that will enable the UK to achieve its dual ambition of AI leadership and ethical excellence.

The paper is structured to provide a comprehensive view of the RAI profession's current state and potential future trajectory:

- **Section 1** maps the current landscape, examining how RAI roles are positioned within organisations and their day-to-day responsibilities
    - *This section establishes the foundation of the paper by revealing how responsible AI roles currently function within organisational structures, demonstrating the operational reality of ethical principles that must be understood before improvements can be suggested.*
- **Section 2** explores practitioner profiles, including diverse entry pathways, organisational structures and core competencies
    - *By exploring the diverse backgrounds and organisational positioning of RAI practitioners, this section highlights both the multidisciplinary strength of the field, its core competencies and the challenges this diversity creates for standardisation efforts.*
- **Section 3** analyses the skills and expertise required for effective practice and educational programmes available to teach these skills
    - *This section identifies current postgraduate and short-term courses and associations available for developing and sustaining a talent pipeline, providing essential insights for hiring managers, educational institutions and professionals seeking to enter or advance in the field.*
- **Section 4** examines the benefits and challenges of professional standards and certification frameworks.
    - *By examining certification frameworks and professional standards, this section addresses the tension between formalising the profession while maintaining accessibility and diversity, a critical consideration for the field's maturation.*
- **Section 5** looks ahead to emerging trends shaping the field's evolution and suggests next steps.
    - *This forward-looking section illuminates emerging trends and evolving responsibilities that will shape RAI practice, helping organisations and practitioners prepare for future challenges in this rapidly developing field.*

Throughout this paper, we balance descriptive analysis with practical insights for organisations seeking to establish or strengthen their RAI capabilities, revealing a field in active formation – one that must balance standardisation and professionalisation with the imperative to maintain diverse perspectives and approaches.

1. Current Landscape

Before we can develop effective pathways into the RAI profession, we must first understand the terrain as it exists today. This section maps the current landscape of RAI practice in the UK – in order to establish an essential foundation for our subsequent analysis and recommendations.

Understanding this landscape matters because effective solutions must be grounded in reality. By documenting the positioning, reporting structure and day-to-day responsibilities of current RAI roles within organisations we can identify both strengths to build upon and gaps that need addressing. This empirical foundation ensures that our recommendations respond to actual challenges rather than theoretical concerns.

Second, the rapid evolution of these roles reflects the increasing strategic importance of AI governance. Tracing this evolution provides valuable insights into how organisations are adapting to growing AI capabilities and regulatory pressures. Mapping the diverse organisational contexts in which RAI practitioners operate reveals why flexible, adaptable frameworks for professional development are necessary. The variation across industries and organisational structures demonstrates that one-size-fits-all approaches to professionalisation would be inadequate.

Finally, by establishing clear definitions of AI, responsible AI and professionalism, we create the conceptual clarity needed to meaningfully discuss how to develop this emerging field. These definitions provide a shared language for the collaborative efforts between industries, professional bodies and policymakers that will be essential for strengthening the RAI ecosystem.

This section, therefore, serves as both foundation and compass for our investigation. It provides a basis for subsequent analysis in current practice and orients our recommendations toward addressing real-world challenges facing the profession.

1.1 Defining AI and Responsible AI

We adopt the OECD definition of AI[5] as a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or

[5] "Home." *AI Principles Overview - OECD.AI*, oecd.ai/en/ai-principles.

virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

Building on this foundation, responsible AI, as defined by ISO,[6] represents an approach to developing and deploying artificial intelligence from both an ethical and legal standpoint. The goal is to employ AI in a safe, trustworthy and ethical way. Using AI responsibly should increase transparency while helping to reduce issues such as AI bias.

1.2 The Foundation of Professionalism and Role of Professional Bodies

This section is dedicated to the question of professional identity and governance. These are fundamental to our investigation because they directly impact how we develop skill pathways, create certification standards and establish credibility in the RAI profession. Without clarity on the appropriate professional context for RAI practice, efforts to build capacity will remain fragmented and potentially ineffective.

The essence of professionalism extends beyond simply working in a particular role. Professionalism at its core is a commitment to the public good. Societies grant professionals significant privileges with the expectation that their work will benefit society as a whole. This social contract lies at the heart of established professions like medicine, law and engineering. We must consider how this fundamental principle of professionalism applies within the emerging responsible AI field.

Professional bodies serve as the guarantors of a profession, upholding its commitment to the public good while acting as arbitrators and enforcers of professional standards. Several established organisations could potentially fulfil this role for RAI, including, but not limited to, ACM, Alliance for Data Science Professionals, TechSkills, BCS and IEEE. The question of what role these incumbent bodies should play in professionalising RAI is crucial for the field's development. The involvement of these organisations or newly created, specialised bodies will significantly shape how RAI professionalism evolves and is governed.

There exist multiple perspectives on how to position RAI within the existing professional landscape. One approach conceptualises AI as a specialised subfield within computer science, suggesting that RAI professionalisation could build upon existing infrastructural and governance models rather than creating entirely new systems. This

[6] "Building a Responsible AI: How to Manage the AI Ethics Debate." *ISO*, 31 Jan. 2024, www.iso.org/artificial-intelligence/responsible-ai-ethics.

approach would offer efficiency and proven methodologies while acknowledging AI's unique challenges.

Alternative perspectives suggest integrating AI governance into established frameworks, such as incorporating it into the UK Cyber Security Council's Career Framework alongside data protection & privacy pathways. This integration would recognise AI's distinct ethical dimensions while providing institutional support through recognised certification bodies.

As the UK seeks to cultivate trust in AI systems through competent professionals, it is essential to determine the right institutional frameworks to support these practitioners in order to achieve broader policy ambitions for responsible innovation at scale. Organisations are increasingly relying on RAI practitioners to navigate complex ethical challenges; thus, there is a pressing need for recognised qualifications or methods to validate the expertise of these practitioners.

1.3 The Evolution of Responsible AI Roles

The current landscape of AI governance is evolving rapidly, and organisations are restructuring themselves in response to advancing AI capabilities, increased AI adoption and evolving regulation and policy regimes.

The proliferation of responsible AI roles represents one of the most significant organisational adaptations to AI becoming mainstream. What began as isolated positions within large technology companies has developed into a recognised professional field across diverse sectors. Ethics and governance now demand an intense strategic focus and extensive change management throughout several industries.

This growth is driven by several key factors:

- **The acceleration of AI capability and adoption** has created a demand for ethical oversight and governance. Organisations are treating ethics as a practical necessity requiring dedicated expertise and resources. This shift is seen across the economy but has been particularly pronounced in heavily regulated industries with low appetites for risk, such as financial services, defence and healthcare.
- **Compliance and regulatory preparation** have become major drivers of role creation and evolution. In the EU, this is driven by organisations actively working towards implementing the EU AI Act and its requirements for high-risk AI systems. Similarly, the UK's pro-innovation approach to AI regulation has created an additional demand for practitioners who can navigate principles-based

frameworks while maintaining an organisation's competitive advantage. Financial services firms are particularly focused on meeting emerging requirements from regulators like the FCA and PRA regarding AI model risk management and governance.

- **The RAI role itself has evolved** from an advisory position to one with direct operational impact on organisations. Early RAI practitioners often served primarily as ethical consultants, but today's roles increasingly involve the direct oversight of AI development and deployment processes, risk assessment and governance implementation.
- **The reporting structure and organisational landscape have shifted** from what began as an auxiliary management within HR and ESG teams or isolated committees. This has evolved into centralised oversight and c-suite management, with the opening of new key positions in organisations such as chief ethics officer. Some practitioners now report directly to senior leadership, present to boards, participate in investor due diligence and work closely with compliance and legal teams.
- **The scope of these roles has expanded** beyond pure ethics to encompass a broader range of responsibilities including regulatory compliance, stakeholder engagement, technical governance and the provision of access to data. This evolution reflects the growing recognition that RAI requires a holistic approach to governance and risk management.

This transformation of RAI roles from peripheral advisors to essential positions substantiates the idea that human expertise is fundamental to effective AI governance. The organisational changes documented here - particularly the elevation of RAI practitioners to positions with direct operational influence - suggest that specialised professional capacity is not merely beneficial but necessary for RAI implementation at scale.

1.4 Current Industry and Sector Context

The demand for RAI expertise varies across industries and organisational contexts, shaped by the interplay of regulatory pressures, organisational maturity and sector-specific requirements. Organisations are responding to these demands in diverse ways, reflecting their operational contexts, risk profiles and existing duties of care. This variation underscores a critical challenge for professionalising RAI practice: any framework for skills development or career progression must be flexible enough to accommodate these diverse contexts while still providing a meaningful structure.

Understanding these sectoral differences is essential for developing targeted interventions that build capacity where it's most needed and in forms that are most appropriate. It also highlights why a one-size-fits-all approach to professional standards or certification would likely be ineffective. By mapping these contextual differences, we lay the groundwork for more nuanced recommendations on how to cultivate RAI expertise across the UK economy in ways that respond to the specific assurance needs of different sectors while maintaining a coherent professional identity. While foundational frameworks exist such as the OECD principles,[7] the UN AI Ethics Framework[8] and the Council of Europe Treaty,[9] their practical implementation varies significantly by region. Organisations are increasingly turning to concrete frameworks like ISO/IEC 42001 (UK),[10] EU AI Act,[11] and the NIST RMF (USA),[12] all of which support the AI Management Essentials Tool from DSIT,[13] a self-assessment tool that aims to help organisations assess and implement responsible AI management systems and processes.

Organisations have approached EU AI Act compliance differently. While some have created dedicated full-time positions, others have distributed these responsibilities across existing roles as additional duties. Data protection professionals are increasingly assuming AI governance responsibilities, as can be seen by an IAPP report,[14] which found that 69% of data protection professionals have incorporated AI oversight into their roles.

---

[7] "Home." *AI Principles Overview - OECD.AI*, oecd.ai/en/ai-principles. Accessed 28 Mar. 2025.

[8] "Principles for the Ethical Use of Artificial Intelligence in The ..." *United Nations System* , High-Level Committee on Programmes (HLCP) Inter-Agency Working Group on Artificial Intelligence, unsceb.org/sites/default/files/2022-09/Principles for the Ethical Use of AI in the UN System_1.pdf.

[9] "Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law." *Council of Europe Treaty Series - No. 225*, 2024, rm.coe.int/1680afae3c.

[10] "ISO/IEC 42001:2023." *ISO*, 18 Dec. 2023, www.iso.org/standard/81230.html.

[11] "The EU Artificial Intelligence Act." *EU Artificial Intelligence Act*, artificialintelligenceact.eu/.

[12] Computer Security Division, Information Technology Laboratory. "NIST Risk Management Framework: CSRC." *CSRC*, csrc.nist.gov/projects/risk-management.

[13] Department for Science, Innovation and Technology. "Ai Management Essentials Tool." *GOV.UK*, GOV.UK, 6 Nov. 2024, www.gov.uk/government/consultations/ai-management-essentials-tool.

[14] Casovan, Ashley, et al. "AI Governance in Practice Report 2024." *AI Governance in Practice Report 2024*, iapp.org/resources/article/ai-governance-in-practice-report/.

The approach to RAI implementation varies significantly based on organisational size and maturity;[15] however, both are inclined to start with AI literacy and responsible guidelines of use policies. Larger organisations and those with more sophisticated AI capabilities typically establish structured, specialised RAI roles with clear mandates and reporting lines afforded by resources. Smaller organisations often take a more integrated approach that involves incorporating the responsibilities of RAI practitioners into existing governance or technical roles. This creates hybrid positions that combine ethical oversight with other functions, allowing organisations to address AI governance needs within resource constraints. The market has responded to this diverse landscape with the emergence of specialised RAI consultancies, assurance firms and an expanding ecosystem of companies offering RAI software solutions as a service (SaaS).

Different sectors have developed distinct approaches to RAI based on their specific risks and requirements. Financial services organisations typically emphasise model governance and bias mitigation. This reflects a focus on fairness and accountability in financial decision-making which is driven by consumer duty regulations that require UK financial services companies to act in the best interest of their customers (fiduciary duties). Healthcare institutions often prioritise patient safety and privacy considerations, developing specialised frameworks for managing AI in clinical settings. Technology companies often focus on ethical product development and scalable oversight mechanisms, while public sector organisations emphasise transparency and accountability to maintain public trust.

The organisational context in which RAI roles operate continues to evolve as companies gain experience with AI implementation and encounter new challenges. Many organisations are transitioning from reactive approaches, where RAI considerations are addressed late in the development cycle, to proactive models that embed ethical considerations from the earliest stages of AI development, ideally before any code is written (ethics by design). This shift reflects the recognition that effective AI governance requires systematic integration into organisational processes that are iterative and run throughout the full lifecycle, rather than implementing interventions only after harm occurs.

Geopolitical shifts in diversity and inclusion policies, along with emerging AI nationalism, are influencing the current RAI landscape and may present new

[15] Dotan, Ravit. "A Flexible Maturity Model for AI Governance Based on the NIST AI Risk Management Framework." *IEEE USA*.

considerations for responsible AI professionals. Organisations implementing RAI frameworks may increasingly navigate complex regulatory environments where certain fairness assessments, such as evaluating models for gender or ethnic bias, become subject to policy debate or face varying regional requirements. These evolving conditions suggest that RAI practitioners may need to develop adaptive strategies for maintaining ethical AI development across different socio-political contexts.

1.5 Day-to-Day Responsibilities of RAI Professionals

The daily work of RAI practitioners involves convening, translating and weighing tradeoffs. At the core of their responsibilities lies governance implementation, where practitioners develop and maintain AI governance frameworks, work to understand and use assurance tools, monitor AI risks, develop AI literacy in organisations and establish reporting mechanisms. This work is increasingly focused on ensuring compliance with emerging regulatory requirements while taking measurable actions to demonstrate alignment with organisational values and objectives.

Stakeholder engagement forms another aspect of the role, with practitioners serving as bridges between technical teams, leadership and external stakeholders. They collaborate closely with development teams to embed ethical considerations into AI design and development processes, while simultaneously advising leadership on AI strategy and updating investors during due diligence processes and when preparing for initial public offering. A significant portion of their time is dedicated to building awareness and understanding of AI ethics across their organisations, often through training programmes or regular consultation with various departments.

Technical oversight can represent another component of RAI practitioner responsibilities. Practitioners and their technical counterparts or team members review AI systems for potential bias, assess algorithmic impact and work to ensure transparency and explainability in AI systems. This involves close collaboration with data scientists and engineers to implement technical solutions to ethical challenges, while maintaining appropriate documentation and monitoring systems.

Strategy and approach to implementing ethics and compliance requires the orchestration of multiple elements. Practitioners draw on leading standards like the EU AI Act, ISO 42001 and NIST RMF, while employing assurance techniques such as algorithmic audits, impact assessments and model cards. RAI practitioners must evaluate this wealth of existing tools - which often differ across jurisdictions – to determine which combinations would be a feasible approach for their organisation. This

involves identifying key decision-makers, locating essential data sources and developing a sustainable strategy for responsible innovation.

However, these practitioners can face significant challenges in executing their responsibilities. Resource constraints often force them to prioritise among competing demands, while the consistent developments of AI technology requires continuous learning and adaptation. One key challenge lies in engaging stakeholders and demonstrating to them the value of RAI practices and how they can support long-term business sustainability and profitability.

The challenge of organisational integration remains persistent, as practitioners must work to establish effective relationships with technical teams and secure sustained buy-in from leadership. This is complicated by the difficulty of defining and measuring success in RAI implementation, particularly when balancing quantitative metrics with qualitative considerations.

RAI practitioners frequently need to develop innovative approaches due to the evolving nature of standards and best practices in this emerging field, which can lead to inconsistencies and constant updating across organisations. The unique characteristics of each AI use case often require tailored ethical frameworks and methodologies, making comprehensive standardisation particularly challenging.

These challenges underscore the ongoing evolution of the field and the need for continued development of professional standards[16] and support structures. As the field matures, addressing these fundamental challenges will be key for establishing effective RAI practice across organisations and ensuring the sustainable growth of this professional domain.

1.6 Team Composition and Dynamics

Successful RAI implementation often depends on effective team composition and collaboration models. Crucially, while RAI practitioners need broaden their understanding across multiple domains, they cannot be expected to possess deep expertise in all areas. Instead, their primary value often lies in their ability to bring together and coordinate diverse experts and stakeholders. High-performing teams succeed because practitioners effectively convene and orchestrate collaboration

[16] Stahl, Bernd. "Living with AI and Emerging Technologies: Meeting Ethical Challenges through Professional Standards." *BCS*, 19 Feb. 2024, www.bcs.org/articles-opinion-and-research/living-with-ai-and-emerging-technologies-meeting-ethical-challenges-through-professional-standards/.

among specialists in areas such as technical development, legal compliance, risk management and data and domain management.

This convening role requires strong relationship-building skills and the ability to facilitate productive dialogue between groups that may have different priorities and perspectives. The professional experts that often must participate in this process include ethicists, legal experts, data scientists, social scientists, domain experts from relevant fields, privacy specialists, policy experts, and the laymen. [17]

Many organisations are moving toward hub-and-spoke models where RAI practitioners act as central coordinators while maintaining networks of subject matter experts across the organisation. This model recognises that effective AI governance requires drawing on distributed expertise rather than centralising all capabilities within a single role or team. The practitioner's success often depends more on their ability to identify when to bring in specific expertise and facilitate effective collaboration than on their individual technical or domain knowledge.

It is essential to understand that the RAI practitioner's role is not to serve as a singular point of accountability for AI-related harms or failures. Rather than creating these roles as a mechanism to deflect or redirect responsibility when issues arise, organisations must maintain clear lines of shared accountability or liability across leadership, development teams and operational units. The practitioner's role is to facilitate and strengthen these accountability structures, not to absorb responsibilities that properly belong at various other levels of the organisation. Effective governance requires commitment and ownership from all stakeholders involved in AI development and deployment.

Establishing a dedicated RAI team within an organisation offers significant advantages over individual practitioners working in isolation: team members can support one another, share diverse perspectives and specialised knowledge and collectively build greater organisational influence. This collaborative approach not only reduces the emotional and intellectual burden on individuals but also creates a more resilient and comprehensive RAI practice capable of addressing the multifaceted challenges of ethical AI development and deployment.

[17] Schuett, Jonas. “How to Design an AI Ethics Board.” *AI and Ethics, Springer*, link.springer.com/article/10.1007/s43681-023-00409-y.

2. Practitioner Profiles

Having established the evolving landscape of RAI implementation and the growing organisational significance of these roles, we now turn our attention to the individuals who occupy these positions. Understanding who RAI practitioners are - their backgrounds, career trajectories, organisational positioning and core competencies - is essential for several reasons.

First, by mapping the diverse pathways into RAI practice, we can identify effective entry points for future talent development and create more accessible routes into the profession. Second, examining organisational structures and reporting lines reveals critical insights on how to position RAI functions for maximum effectiveness and influence. Third, analysing the core competencies of successful practitioners provides the foundation for developing targeted training programs and professional standards that build the necessary capabilities.

This section addresses a central question: how to strengthen the UK's AI assurance ecosystem as called for in the AI Opportunities Action Plan. To support the human infrastructure, it relies on, that is, the RAI practitioner, we must first understand who these professionals are, how they operate within organisations and what skills make them effective. Only with this understanding can we design interventions that build on existing strengths while addressing gaps in the professional ecosystem.

2.1 Career Pathways and Backgrounds

The pathways leading to RAI practice are remarkably diverse, reflecting the multidisciplinary nature of the field. RAI practitioners come from varied backgrounds including philosophy, compliance, computer science, law and social sciences and business management. Recognising these diverse pathways into RAI practice is critical for developing inclusive talent pipelines. By understanding how current practitioners entered the field, we can identify multiple entry points for future professionals, ensuring that the field continues to benefit from diverse perspectives while creating more structured routes for those seeking to join the profession.

The diversity of backgrounds can also be seen as reflective of the relative immaturity of RAI as a professional discipline. Its current state can be compared to privacy practice 20 years ago, where defined career paths had not yet emerged and practitioners entered the field from different directions, based on their existing expertise and organisational needs. As the profession matures, we can expect more standardised educational and career pathways to develop. However, the inherently multidisciplinary nature of RAI

work suggests that maintaining diversity in professional backgrounds will remain valuable even as the field evolves.

Many current RAI practitioners describe their journey as non-linear, often beginning in traditional roles before transitioning into AI ethics as the field evolved. Early practitioners frequently report creating their positions organically, identifying ethical gaps in AI development processes and gradually expanding their scope to address these challenges. Those with technical backgrounds often cite moments of ethical awakening that led them to focus on responsible implementation, while those from humanities or social science backgrounds describe intensive self-directed learning to build competencies.

One notable trend is the increasing number of practitioners entering the field through dedicated academic programmes in AI ethics and governance. These programmes, while relatively new, are beginning to create more structured pathways into the profession (see section 3 for the current educational landscape). However, the majority of current practitioners still emphasise the importance of continuous learning and adaptation, regardless of their initial background.

2.2 Organisational Structure and Reporting Lines

The positioning of RAI functions within organisational hierarchies directly impacts their access and influence. By examining various structural approaches, we can identify models that maximise impact while maintaining the necessary independence of RAI practitioners. This analysis provides essential guidance for organisations seeking to establish or strengthen their RAI capabilities.

Common reporting structures include direct lines to chief technology officers, chief data officers, or, increasingly, dedicated heads of AI governance, while working closely with compliance teams or data offices. Some organisations have established independent AI ethics committees or boards, with practitioners serving as key liaisons between these oversight bodies and operational teams.

We have seen an evolution in organisational structures, with many companies moving from distributed responsibility models to more centralised approaches. This shift could reflect the recognition that effective AI governance requires dedicated resources and clear lines of accountability. However, successful implementation frequently depends on maintaining strong collaborative networks across departments while ensuring independence in ethical decision-making.

Engagement with executive leadership remains a key factor in effectiveness. Practitioners with direct access to senior leadership report great success in implementing ethical frameworks and driving organisational change from the top down. This access becomes particularly essential when addressing sensitive issues or recommending significant changes to AI development processes.

2.3 Core Competencies

Defining the essential skills and knowledge that enable effective RAI practice forms the foundation for professional development frameworks. This analysis helps address a key challenge identified in our introduction: the need for clear skill pathways that hiring managers can recognise and aspiring practitioners can develop. By articulating these competencies, we hope to support a critical step toward establishing recognised professional development pathways.

Through existing literature[18] and discussions with current practitioners, several fundamental competencies have emerged as essential for effective RAI practice. Technical and data literacy, while not always requiring deep expertise, proves helpful for engaging meaningfully with development teams and understanding AI systems' potential impacts. Practitioners must be able to understand technical documentation, engage in discussions about system design and identify potential ethical issues in technical implementations.

Equally important are strong analytical and problem-solving skills, particularly in navigating complex ethical dilemmas that lack clear solutions. Successful practitioners demonstrate the ability to balance competing interests, acknowledge existing incentive structures, analyse potential impacts across diverse stakeholder groups and develop practical solutions that align with both ethical principles and business objectives.

Robust analysis and research capabilities prove useful, as practitioners must excel at identifying and evaluating ethical uncertainties and value conflicts across diverse consultation scenarios. This requires the ability to effectively navigate and apply relevant ethics literature, policies, guidelines and standards to inform decision-making processes.

[18] Cocchiaro, Mariangela Zoe, et al. "Who Is an AI Ethicist? An Empirical Study of Expertise, Skills, and Profiles to Build a Competency Framework." *SSRN*, 17 July 2024, papers.ssrn.com/sol3/papers.cfm?abstract_id=4891907.

Process management and strategy implementation form the practical backbone of the role. Practitioners assess stakeholder impact and quantify risk magnitude while guiding the implementation of chosen solutions.

RAI practitioners lead formal consultations and stakeholder meetings and maintain comprehensive documentation of consultation processes and outcomes. Critical to this work is their ability to identify systemic issues and understand their broader implications. Complementing this approach, knowledge of regulations, governance frameworks and industry standards strengthens the practical application of ethics in real-world scenarios.

Interpersonal and educational leadership skills complete the practitioner's toolkit. Success in the field demands active listening and empathy in stakeholder engagement and fostering effective cross-functional communication and collaboration. These skills are often enhanced by a foundation in philosophical study that enriches ethical reasoning and debate. Practitioners must excel at translating complex ethical concepts into compelling narratives and practical frameworks that resonate across diverse organisational contexts. RAI practitioners must have an ability to build consensus around ethical approaches and influence organisational to achieve success in the role. They may develop and deliver ethics training programmes for organisational staff while advocating for stakeholder interests and maintaining objectivity. Drawing on both philosophical foundations and storytelling capabilities, they navigate ethical debates with a balanced perspective, ultimately cultivating an environment conducive to ethical reasoning and discussion.

3. Cultivating the Talent Pipeline – Educational Opportunities

Having established who RAI practitioners are and how they are positioned within organisations, we now turn to the critical question of how to develop and sustain the professional capabilities needed for effective AI governance.

The UK's ambition to be an innovator in AI depends directly on building human capacity to implement governance frameworks. As noted in recommendations 14, 15, 17 and 19 of the AI Opportunities Action Plan, developing talent with the right mix of technical and ethical skills is essential for maintaining the UK's competitive edge while ensuring that AI systems remain aligned with societal values. This section examines how we can systematically cultivate these capabilities at scale by providing suitable training for a talent pipeline while considering the current educational landscape.

3.1 Current Postgraduate Courses

The rapidly evolving nature of AI technology and governance creates continuous learning demands for RAI practitioners. Current training approaches combine formal education with extensive on-the-job learning and professional development. Organisations are increasingly investing in structured training programmes that cover both technical and ethical aspects of AI governance; many practitioners rely heavily on self-directed learning to stay current with developments in the field.

The educational landscape is evolving rapidly to meet this growing demand. Specialised postgraduate degrees in AI ethics, governance and responsible innovation are emerging at leading institutions in the UK whether in the form of master's degrees or summer programming. These programmes typically combine ethical frameworks, governance principles, technical fundamentals and practical implementation strategies.

Please note that we do not endorse, promote, or verify any specific courses listed in the tables below. The following examples are presented solely to illustrate the types of educational offerings currently available in the UK's landscape. As with all lists, some offerings may inadvertently be omitted.

Title	Focus	Length	Location
University of Cambridge, MA in AI Ethics and Society	This program develops leaders capable of addressing critical AI challenges in professional settings, including privacy, surveillance, justice, fairness, algorithmic bias, misinformation, Big Data and responsible innovation practices.	22 months	Cambridge
Brunel University, MSc Human-AI Interaction	This programme bridges the gap between the boardroom and AI tech developers. It is an interdisciplinary non-technical, business application and adoption programme designed and delivered by an interdisciplinary team. The programme covers five core topics: AI technologies for decision making; digital strategies and transformation; ethics, power and inclusion, business models and leadership; and regulatory and legal frameworks. The programme takes students into companies to apply their learnings.	12 months	London
University of Edinburgh, MSc in Data and AI Ethics	This program addresses the growing demand for interdisciplinary expertise in ethical design, implementation and governance of AI and data-intensive technologies, equipping graduates to help organisations navigate complex technical systems with transparency, accountability, fairness and respect for human rights.	12 months	Edinburgh
Northeastern University London, MA in Philosophy and AI	This program examines the ethical and theoretical challenges posed by AI's increasing presence across modern society, providing students with critical thinking, communication and technical skills sought after in professional services, creative industries and government sectors.	12 months	London
Birkbeck University, MA in AI,	This program offers a deep understanding of the relationship between AI technologies and societal values, focusing on AI decision-making processes, fairness, biases, security	12 months	London

Ethics and Society	risks and privacy concerns, while exploring how these technologies impact politics, personal relationships and societal structures across various domains.		
University of Oxford, MSc in Social Science of the Internet	The MSc in Social Science of the Internet equips students with knowledge and skills to critically examine how the internet and emerging digital technologies impact our lives, societies, businesses and governments. Students learn from the Oxford Internet Institute's world-leading multidisciplinary faculty, gaining expertise in key social science concepts, theories, research methods and technological fundamentals. Graduates develop the ability to conduct and communicate high-quality research across the digital domain.	10 months	Oxford

Traditional MSc/MA degrees in AI ethics and governance represent a substantial academic commitment, typically spanning 12-24 months of full-time study and requiring 120-180 credits of coursework. These programs offer comprehensive theoretical foundations, research opportunities and academic credentials recognised globally, culminating in a dissertation that demonstrates mastery of the subject. One might also see ethics modules integrated into technical computer science degrees. We must continue to support a pipeline of industry-funded AI master's courses through the appropriate intermediaries to address the high demand for specialised AI skills while prioritising diversity, inclusion, ethics and competence among graduates.

3.2 Current Short-term Courses

In contrast, shorter-term offerings such as professional certificates and specialised online courses provide targeted knowledge acquisition in specific aspects of AI ethics and governance, usually completed in 12-30 hours over weeks rather than years. While these accelerated options offer flexibility, immediate practical application and lower financial barriers, they necessarily sacrifice the breadth, depth and academic recognition of university degrees.

The choice between these educational pathways depends largely on career goals, time availability, financial resources and whether one seeks comprehensive expertise or focused skill enhancement in particular aspects of responsible AI.

Alongside traditional education we see a proliferation of online dedicated courses. These include, but are not limited to the following:

Title	Focus	Length	Location
IEEE CertifAIEd Authorised Assessor	A certification program evaluating ethical standards of autonomous intelligent systems (AIS), providing certification guidance, assessment and independent verification to scale responsible innovation and improve AIS quality while building stakeholder confidence.	28 hours	Online
IAPP, Certified AI Governance Professional	Training programme developing professionals with expertise in AI governance, risk reduction, compliance improvement and the implementation of appropriate controls and procedures to enhance brand loyalty.	13 hours	In person and online options
BlueDot Impact, AI Safety Fundamentals	An intensive five-day course exploring transformative AI's technical foundations and impact on humanity through expert-facilitated discussions, examining potential futures and key debates alongside peers passionate about ensuring that AI benefits humanity.	15 hours	Online
The Chartered Institute for Securities and Investment's Certification in Ethical AI	Professional assessment covering AI foundations, ethical considerations, global regulatory developments and risk management strategies, with an emphasis on strategic and tactical implementation for executives.	12 hours	Online
ForHumanity courses and certifications	Range of certifications from the ForHumanity Certified Auditor (FHCA) to introductory certifications for executives seeking general knowledge and professionals building foundations of expertise in AI risk management.	14 hours	Online
LocialAI Courses	Interactive online sessions designed for professionals seeking to enhance expertise in AI governance through flexible, immersive learning focused on practical implementation skills.	6 hours	Online

BABL's AI and Algorithm Auditor Certification Programme	Training programme for AI governance, risk and compliance professionals, providing knowledge in AI risk assessment, governance best practices and common techniques in modern AI/ML development.	30 hours	Online
BCS Foundation Certificate in the Ethical Build of AI	This course introduces the key issues in ethical AI and equips learners to design and build AI ethically, navigating risks from security to human rights. This course is ideal for anyone involved with designing or developing software that uses AI.	12 hours	Online
HITRUST Academy's AI Risk Management Certification	HITRUST has developed a dedicated AI Security Risk Management and Certification Framework, supporting the secure development and deployment of AI systems. The AI Security Assessment & Certification Specification is informed by a multi-stakeholder working group including AI developers, cybersecurity experts, regulators, and healthcare representatives.	20 hours	Online

While these existing specialised certifications and courses can establish recognised standards of professional competency and provide structured pathways for skills development, professionals and students need guidance on selecting programmes that are high-quality, industry-relevant and recognised in the UK.

Alongside courses, mentorship and peer learning are central development mechanisms in this field. While theoretical knowledge is essential, experienced practitioners emphasise that it must be complemented by exposure to real-world cases and challenges. The development of sound judgment and decision-making capabilities requires guidance through complex scenarios from seasoned professionals.

This need for practical wisdom has fostered vibrant communities of practice, both online and offline, where RAI practitioners actively share insights, resources and support. These range from informal peer networks (such as All Tech is Human and Responsible AI UK) to established associations like the International Association of

Privacy Professionals, the Association of AI Ethicists and the International Association of Safe and Ethical AI, creating a collaborative ecosystem where practitioners at all levels can learn from others' experiences and challenges. Alongside leading consortiums and global communities such as AI Verify Foundation's Project Moonshot, Partnership on AI, and The AIGI Consortium. These communities function as 'professional incubators', establishing norms and practices that may eventually evolve into more formal standards. Through case sharing, collaborative problem-solving and the development of shared assessment methodologies, these networks are helping to bring consistency to RAI implementation across organisations.

These developments represent important steps toward professionalisation while highlighting the need for coordination and the alignment of standards across different certification bodies. As the field matures, greater convergence around core competencies and assessment methods will likely emerge, though maintaining flexibility and relevance across different organisational contexts remains essential. Adaptability remains essential - particularly in the UK context, where a sector-specific approach to AI governance may require the flexible application of these foundational skills across different industries.

4. Standards and Certifications: Balancing Rigour, Accessibility and Ethics

As the RAI profession continues to evolve from an emerging field into an established discipline, organisations, practitioners and policymakers face a critical decision point: should this work be formalised through professional standards and certification frameworks, and if so, how is this to be done? This question directly connects to our central investigation of developing clear pathways into the profession and supporting the UK's RAI practitioner ecosystem and, in turn, the assurance ecosystem.

The UK's AI Opportunities Action Plan explicitly calls for the development of a robust AI assurance ecosystem (Recommendation 29) and the strengthening of the UK's skills pipeline (Recommendations 14-19). Professional standards and certifications represent potential mechanisms for achieving these goals by establishing recognised benchmarks for practitioner competence, creating structured development pathways and building credibility for the profession. However, these same mechanisms could potentially undermine the field's diversity and accessibility if implemented without careful consideration.

This section examines the potential benefits, drawbacks and implementation challenges of formalising RAI practice through standards and certifications. Understanding these trade-offs is essential for making informed decisions about professionalising RAI in ways that strengthen rather than constrain the UK's capacity for responsible innovation. We hope that the section that follows will help organisations, educational institutions and policymakers navigate the complex choices involved in developing professional frameworks that balance rigour with inclusivity.

By ‘certification options’, we refer to the various approaches that could be used to formally recognise and validate RAI practitioners' expertise. These range from academic credentials (such as specialised degrees and certificates) to industry-recognised professional certifications, competency frameworks and continuing professional development requirements. Each approach represents a different model for establishing professional credibility and ensuring practitioner competence.

4.1 Benefits of Formalising the RAI profession

The movement toward formalising RAI practice through standards and certifications represents a natural evolution of the field. As we have seen in other industries, formalisation brings much-needed clarity to role definitions and career pathways. This formalisation could help organisations better structure their AI governance functions.

For instance, the UK's financial services sector has built its global reputation partly on the depth and quality of its risk management and compliance practitioner skills. Similarly, the UK's AI industry has an opportunity to develop a competitive advantage through the strength of its RAI practitioner community. This parallel suggests that investing in RAI talent development could become a cornerstone of the UK's AI strategy, creating a skilled workforce that supports responsible innovation while maintaining international credibility.

Professionalisation of the RAI field not only boosts the credibility of practitioners, but also provides greater clarity around career progression and talent development. Certification can signal both rigour and trustworthiness—enhancing the legitimacy of ethical roles in AI development and helping organisations reduce operational and reputational risk. As AI systems become more embedded in regulated domains, recognised certifications can also provide reassurance to regulators and the public alike.

Standardisation can further support recruitment and professional development by establishing clear benchmarks for skills and expertise, enabling the creation of a robust professional ecosystem similar to that seen in financial services. Formal certifications and standards help build credibility for and access to the profession, providing organisations with confidence in practitioners' capabilities. This approach increases employee buy-in during implementation (as colleagues may be more likely to sincerely adopt recommendations). Standardisation can also accelerate the development of training programmes and educational curricula, creating clearer pathways for talented individuals aspiring to enter the field.

Currently, a significant challenge facing the profession is the lack of clearly defined role boundaries and responsibilities. Formalisingthe profession may assist with this. The RAI practitioner's title is in danger of becoming a catch-all moniker for professionals who deal with various AI-related challenges, from technical governance to stakeholder communications, policy monitoring and risk assessment. This expanding scope of responsibility, while reflecting the role's importance, can lead to practitioner burnout and diluted effectiveness. Organisations should carefully balance the need for comprehensive AI oversight with realistic expectations of what one role or team can achieve.

4.2 Challenges of Formalising the RAI profession

While standardisation offers clear benefits, it also presents significant challenges. Overly rigid certification requirements risk excluding valuable practitioners with non-

traditional backgrounds. This could inhibit the field's diversity and close off vital, interdisciplinary routes into the profession. For many current RAI professionals, particularly those with backgrounds outside of engineering or computer science, certification must remain accessible and flexible.

Furthermore, there is growing concern that formal ethics work may drift into a compliance-only function—reducing complex, context-sensitive ethical deliberations to procedural checklists. The challenge for professionalisation will be to ensure that ethical work maintains its independent, deliberative character rather than becoming subsumed into compliance regimes.

One challenge lies in maintaining distinct professional streams within AI governance. Ethics work risks devolving into regulatory checklists and losing crucial sociotechnical perspectives. While certification programs may address RAI deployment protocols, they often fail to cultivate the philosophical reasoning, logic and creative problem-solving needed for ethical AI development, especially when working with sensitive data and vulnerable communities.

It is crucial to recognise that academic AI ethics, which focuses on theoretical frameworks and foundational research, serves a different purpose from practical RAI implementation in organisational contexts. While these streams should inform each other, they require different skill sets and serve distinct purposes. The boundaries between ethics, audit and compliance roles should be clearly defined so as to maintain the integrity and effectiveness of each function.

The question of certification authority is equally challenging. Commercial bodies may prioritise profit over accessibility, academic institutions might create prohibitively expensive programs and industry associations, while understanding practical needs, may lack broader societal perspectives. These concerns raise important questions about representation in shaping the profession's future.

Economic implications require careful consideration. Existing practitioners, particularly those from underrepresented groups, could face pressure to invest in costly recertification. Without deliberate effort to address accessibility barriers, especially those faced by practitioners from lower socioeconomic backgrounds, practitioners with disabilities, or caregivers requiring flexible options, certification requirements could entrench existing inequities. These risks could create a two-tiered system where certification becomes a privilege of those who can afford it rather than a true measure of competency.

These challenges reflect broader patterns in professional certification systems, which typically balance restricted supply (through rigourous requirements) with quality assurance and economic sustainability. The tension between accessibility and maintaining high standards requires thoughtful navigation.

Professional bodies typically operate on a model where practitioners command higher fees for their services, which, in turn, finance the costs of certification and ongoing professional membership. While the substantial upfront investment required for certification (whether in law, medicine, engineering, or other fields) can create barriers to entry, this system serves multiple purposes beyond simple access.

The restricted supply of certified professionals, maintained through rigourous qualification requirements, helps ensure that practitioners can recoup their initial investments through high service fees. This economic model also supports the quality of professional education, as can be seen by the comprehensive training programmes required for various certifications. While some might view this as exclusionary, there is an argument that demanding qualification standards serve to ensure competency in fields where expertise is crucial for public safety and effective service delivery.

This creates a complex tension between accessibility and professional standards. Making certification more affordable or eliminating cost barriers could increase the supply of professionals and potentially reduce service costs to the public. However, this might also reduce the incentive to maintain rigourous standards and could undermine public confidence in the certification's value. The challenge lies in balancing these competing interests while addressing valid concerns about equity and access.

In this context, new certification requirements in this field must be evaluated not only on their immediate impact on existing practitioners but also on how they fit into this broader pattern of professional qualification systems. The concerns about creating barriers for underrepresented groups remain valid, and addressing this problem may require engaging with more fundamental questions about how professional certification systems can be structured to balance accessibility, economic sustainability and the maintenance of high professional standards.

The rapid evolution of AI technology poses additional challenges, as certification frameworks must remain flexible enough to accommodate emerging ethical issues. The scope of certification also requires consideration—should it focus specifically on AI ethics or encompass digital ethics more broadly?

Ultimately, the key challenge is preserving the field's inherent diversity and interdisciplinary nature while providing meaninguful structure. Successful approaches will need to recognise multiple paths to expertise and create flexible frameworks that acknowledge different specialisations, from governance to stakeholder engagement. The goal should be to create professional standards that enhance rather than constrain the field's ability to address the multifaceted challenges of ethical AI implementation.

5. Future Outlook and Conclusions

As we conclude our mapping of the RAI profession's current state, we must look ahead to the challenges and opportunities that will shape its evolution. This forward-looking analysis is essential for several reasons. First, the rapid pace of AI innovation means that the responsibilities and required competencies of RAI practitioners are continuously evolving. Organisations and practitioners need to anticipate these changes to maintain effective governance as AI capabilities advance. Second, the UK's regulatory approach - which emphasises principles and sector-led governance - places particular importance on having skilled professionals who can operationalise ethical principles without prescriptive regulatory frameworks. Third, as the UK positions itself as a global AI leader, the development of robust RAI professional capacity represents a potential competitive advantage, enabling organisations to demonstrate trustworthiness.

This section examines emerging trends that will shape RAI practice in the coming years and presents actionable recommendations for key stakeholders - organisations, professional bodies and policymakers - who all play essential roles in nurturing this critical professional community.

5.1 Evolving Role Expectations and Emerging Responsibilities

Understanding how RAI roles will evolve is critical for both current practitioners and organisations planning their AI governance approaches. As AI systems become more powerful, autonomous and embedded into critical functions, the scope and complexity of ethical oversight will increase accordingly. By identifying emerging responsibilities now, we can begin developing the competencies, frameworks and support structures needed to address tomorrow's governance challenges before they become critical. This foresight is particularly important in the UK context, where sector-specific regulation places greater responsibility on practitioners to navigate the landscape to ensure responsible innovation. The UK’s regulatory approach makes the role of the RAI practitioners even more critical for maintaining consumer, citizen and industry trust.

The continuous evolution of regulatory requirements and assurance techniques adds another layer of complexity to this process. Practitioners must constantly update their knowledge and adapt their approaches, often while maintaining their existing responsibilities. Organisations are increasingly seeking practitioners who can bridge the gap between regulatory requirements, ethical principles and practical implementation. However, this breadth of responsibility must be matched with appropriate resources,

support structures and clear role definitions to ensure sustainable and effective practice.

The scope of RAI practice continues to expand alongside new AI capabilities. With continual technical developments consistantly generating new ethical implications, this role is one of lifetime learning. For example, small AI systems that can run on edge devices with limited computational resources, multimodal AI systems that combine different types of data and outputs (text, image, audio, video) and agentic AI systems that can autonomously plan and execute complex tasks. Each of these developments brings unique governance challenges and requires new approaches to ethical oversight and risk management.

The role is evolving to include the oversight of AI supply chains, particularly as organisations adopt third-party AI services and models. This becomes more complex when dealing with systems that combine multiple AI capabilities or operate autonomously. Practitioners must develop frameworks for assessing and monitoring these increasingly sophisticated systems while ensuring their responsible deployment across different computational environments.

Environmental and energy considerations are becoming a part of RAI practice, with growing emphasis on sustainable AI development. Practitioners must measure AI's environmental impact alongside its social and ethical implications - evaluating energy consumption, computational efficiency and broader environmental consequences. This includes assessing trade-offs between edge deployment, specialised models and centralised computing. A core challenge remains securing appropriate data access to measure, evaluate and communicate environmental footprint.

We anticipate growing specialisation within the field as practitioners develop expertise in sector-specific domains such as healthcare AI governance, financial services AI ethics and public sector AI accountability - a trend that will become increasingly important as industries develop tailored approaches. While this specialised knowledge is valuable, a foundation of shared principles and methodologies must continue to underpin all RAI practice. In the future, we must deliberately prevent the formation of isolated professional silos by fostering collaborative initiatives that facilitate cross-industry knowledge exchange and maintain cohesion within the broader RAI community.

Finally, the trend toward SaaS for governance tools and metrics is likely to accelerate, though practitioners will need to ensure these tools support rather than replace human judgment in decision-making. The development of the AI assurance ecosystem, as

anticipated by DSIT, will create new opportunities and requirements for RAI practitioners to engage with third-party assessors and auditors.

5.2 Next Steps

Our analysis of the RAI profession reveals both its importance to the UK's AI ambitions and the challenges it faces in continuing to mature as a field. The recommendations presented here directly address the gaps identified throughout our mapping of the profession - from unclear career pathways to insufficient organisational positioning and underdeveloped professional frameworks.

By taking concrete actions now, stakeholders can strengthen this essential professional community before AI governance challenges outpace our capacity to effectively address them. The specific priority actions for each stakeholder group provides ways in which we can cultivate the human infrastructure needed to ensure that AI development in the UK remains innovative and responsible.

Our investigation into the UK's responsible AI profession reveals a critical workforce developing at the intersection of ethics, technology and governance. These practitioners from diverse professional backgrounds serve as essential bridge-builders who operationalise ethical principles and regulatory requirements within organisations.

The profession stands at a pivotal development stage, evolving from advisory roles to strategic functions with direct influence on AI development. Without these professionals to implement principles of safety, transparency, fairness, accountability and contestability, the UK's regulatory approach risks remaining only a conceptual aspiration rather than becoming a practical and operational system.

Priority Actions for Organisations: Establish RAI roles with clear mandates and sufficient authority to influence AI development proactively. Invest equally in technical capabilities and governance skills when developing AI talent. Ensure that RAI practitioners have direct reporting lines to senior leadership. The following priorities represent the most urgent actions needed to strengthen this crucial professional community.

Priority Actions for Professional Bodies: Develop flexible certification frameworks that recognise multiple pathways to expertise. Centre current practitioners in professionalisation discussions to build upon existing best practices. Create accessible professional development opportunities that maintain diversity while establishing standards. Define clear boundaries between the ethical, auditorial and compliance functions of RAI practice. Ensure that emerging certification frameworks accommodate

a wide range of entry routes and validate both formal and experiential learning, especially in ethics, social impact, and interdisciplinary practice.

Priority Actions for Policymakers: Recognise RAI practitioners as essential human infrastructure for effective AI governance, adoption across the economy and development of the assurance ecosystem. Support industry collaboration through networks like techUK to address common challenges. Invest in educational pathways and talent pipelines that develop both technical and ethical competencies. Monitor the profession's evolution to identify areas requiring additional support.

For the UK to achieve its ambition increase the adoption of the AI and develop the AI assurance ecosystem, we must move beyond asking whether organisations need RAI expertise and focus instead on how to effectively develop, deploy and support these professionals across the economy.

techUK
FOR WHAT COMES NEXT


linkedin.com/company/techuk


youtube.com/user/techUKViews


@techuk.bsky.social


info@techuk.org

Image credit | iStock by Getty Images